

Modeling emergence of norms in multi-agent systems by applying tipping points ideas

Francisco Lopez
 School of Computer Science
 University of Manchester
 Manchester, United Kingdom
 lopezfrancisco547@gmail.com

ABSTRACT

Norms are known to be a major factor determining humans behavior. It's also shown that norms can be quite effective tool for building agent-based societies. Various normative architectures have been proposed for designing normative multi-agent systems (NorMAS). Due to human nature of the concept norms, many of these architectures are built based on theories in social sciences. Tipping point theory, as is briefly discussed in this paper, seems to have a great potential to be used for designing normative architectures. This theory deals with the factors that affect social epidemics that arise in human societies. In this paper, we try to apply the main concepts of this theory to agent-based normative architectures. We show several ways to implement these concepts, and study their effects in an agent-based normative scenario.

Keywords

norms, multi-agent systems, tipping points

1. INTRODUCTION

Human societies are simultaneously frustratingly unchanging and yet susceptible to “epidemics” that sweep across the social fabric causing people to adopt previously rare practices. Tipping point theories attempt to explain the subtle triggers behind these social processes. In 2000, Malcolm Gladwell [25] produced a popular science book summarizing three key factors which trigger tipping points: 1) scale-free networks (the Law of the Few); 2) effective messaging (the Stickiness Factor) and 3) environmental influences (the Power of Context). This paper relates tipping point theory to the process of norm emergence in multi-agent systems; we propose that normative agent architectures can serve an excellent computational model for expressing many contagious social phenomena, including tipping points and information cascades.

Social norms are known to be a major factor governing hu-

mans' behavior; unbeknownst to us, many of our everyday behaviors are influenced by these implicit standards [10]. Various normative architectures have been proposed for designing normative multi-agent systems (NorMAS) capable of reasoning about norm adoption. Some of these systems have been grounded in social science theory, but the aim of many architectures is simply to effectively address standard multi-agent system challenges, including agreement formation, coordination and conflict resolution [9].

Despite recent research progress in the area, the complete life-cycle of norms is far from fully understood. The complex nature of human decision-making makes comprehending the rationale behind social interactions difficult, since people are notoriously bad at self-reporting their motivations [15]. The field of agent-based modeling aims to create agents in the image of humans. These agents typically have cognitively-inspired decision-making components, and are situated in life-like scenarios. In both standard multi-agent systems and cognitively-inspired models, existing social theories have been employed toward the construction of normative models [5]. Various stages of the norm life-cycle including recognition, adoption, compliance and emergence are often modeled on similar concepts in social sciences.

This paper proposes a unified model of how norm emergence in networked agent societies can be used to predict the effects of common tipping point triggers. Previous work on norm emergence in networks has investigated the effects of social network topology in static [34, 31] and also dynamic networks [29]. Yu et al. [39] presented an evaluation of different learning methods on norm emergence in networked systems. In our work, we simply employ network structures as a medium to apply ideas from tipping point theory relating to the Law of the Few. Therefore, the structure of agents' network is not of interest by itself, other than making it congruent with human social networks. Our main contribution is showing the role and significance of tipping point principles in normative agent systems, and evaluating the potential impact of this model on NorMAS design. The next section provides an overview of related work in this area.

2. RELATED WORK

Hollander and Wu [24] refer to three categories of normative studies in the social sciences: 1) social function of norms [15], 2) impact of social norms [7], 3) mechanisms leading to the emergence and creation of norms [14]. In the con-

text of social function, norms are often concerned with the *oughtness* and *expectations* of agent behavior; where oughtness refers to the condition where an agent should or should not perform an action, and expectation refers to the behavior that other agents expect to observe from that agent [13, 2]. An example of work belonging to this category is Boella and Torre’s architecture containing separate subsystems for counts-as conditionals, conditional obligations, and conditional permissions [16].

Within the second category, social impact, norms are considered in terms of cost provided to or imposed on the parties involved in a social interaction [4]. For instance, punishment and sanctions are introduced as two enforcement mechanisms used to achieve the necessary social control required to impose social norms [33]. Here they demonstrate a normative agent that can punish and sanction defectors and also dynamically choose the right amount of punishment and sanction to impose [12].

As noted in [24], the third category is concerned with the *how* of norms more than the *why*. Relevant work in NormAS domain that falls into this category is divided into the following groups [11]. The first group is composed of research in which norms are practically hard-wired into the system or dictated directly to the agents. In the second set, norms emerge from social interactions among agents. Sen and Airiau’s work [32] in which agent interactions are modeled using payoff matrices, focuses on norm emergence through social learning in agent societies. Conte et al.’s *EMIL* was designed as an architecture for modeling norm emergence [17]. The EMIL architecture includes a dynamic cognitive model of norm emergence and innovation [6].

Much of the existing work in normative multi-agent systems explicitly or implicitly relies on social science theories [8]. In a recent work, some of the well-known theories of philosopher David Hume were evaluated using an agent-based model called HUME_{2.0} [17]. This work demonstrates how social justice concepts can even emerge from heterogeneous agents that are not endowed with norm representations.

Self-determination theory is also referenced by some of the normative works [3]. Here the focus is on the agents’ motivation and the extent to which the motivation is intrinsic or extrinsic. Neumann studied existing normative architectures to see how much they comply with self-determination theory [27].

Practice theory is an example drawn from anthropology; this theory describes how changes in the society are based on the interactions between the human agents and social structure. For instance, an agent-based model for energy demand and supply social practices is presented in [1], which shows how energy consumption norms form and evolve in urban societies. The next section provides background on tipping point theory, primarily from a sociology perspective.

3. TIPPING POINT THEORY

The term, “tipping point” was initially coined in physics to describe the situation in which the state of an object rapidly changes from one stable equilibrium to another different

equilibrium. Morton Grodzins was the first to use term in social sciences to describe an interesting phenomenon he observed in the US cities, known as *white flight* [23]. His observation was that in some metropolitan areas the percentage of African-American people would increase up to a certain point. After that point, those with white ethnicity immigrated from those cities in large numbers. Thomas Schelling presented the *general theory of tipping*, which describes how individuals’ micromotives and microbehavior can aggregate in the big picture [30]. The model of collective behavior introduced by Mark Granovetter [22] uses thresholds to determine the path of social events. This model was initially used to describe how fads are created.

In normative studies, tipping points are usually denoted as the point of maximum return at which time the behavior has the highest level of acceptability from the population. For instance, in a certain group of friends, the number of times they shower in a week may vary, but a specific value has the highest acceptability by group members as the conventional pattern of behavior. In this paper, we study the impact of Gladwell’s three factors on norm emergence in agent-based normative systems and demonstrate practical ways to apply this versatile theory.

4. EXPERIMENTAL SETUP

For our experiments, we employ the classic scenario, rules of the road, that is frequently used to study normative behavior in multi-agent systems. In this scenario, there exists a population of agents that do not have any preference toward driving on the left or right side of a two-way road. No rules or higher enforcement exist to determine the preferred side. This scenario represents a two-action stage game that models the situation where agents need to agree on one of several equally desirable alternatives. The societal norms that we would like to evolve are either driving on the left or driving on the right [32].

In this scenario agents receive a fixed value reward and punishment based on the following payoff matrix shown in Table 1.

	left	right
left	1,1	-1,-1
right	-1,-1	1,1

Table 1: Payoff matrix for rules of the road scenario

As Yu et al. [39] note, although this payoff matrix appears simple, the coordination game poses a very challenging puzzle for human beings to solve efficiently. The game has two pure Nash-equilibria: both agents drive left or both agents drive right. Classical game theory, however, does not give a coherent account of how people would play a game like this. The conundrum is that there is nothing in the structure of the game itself that allows the players (even purely rational players) to infer what they ought to do. In reality, people can play such games because they can rely on some contextual cues to agree on a particular equilibrium [38].

In similar studies on normative systems, usually the cumulative payoff (reward) of the whole population of agents is used as a measure of comparing various methods (see [32] and [39]

for examples). Instead, we opt to use the norm emergence time for each method as an evaluation metric. This is functionally equivalent since the payoff received by all agents post norm emergence is the same, hence a method which leads to faster norm emergence will also yield the higher cumulative payoff.

5. KEY FEW MEMBERS

In this section, we study the effects of key members of an agent society on the rate of norm emergence. These key members are selected using standard heuristics for measuring influence within a network; we evaluate the performance of three centrality measures: degree, closeness, and betweenness. Degree centrality measures the number of edges connected to a node. Closeness is calculated based on the total distance to all other nodes. Nodes with a high betweenness centrality fall on a large proportion of the shortest paths (geodesics) in the graph.

To model the characteristics of a real social network, we use an algorithm introduced in [35] to create a synthetic network which follows power law degree distribution and exhibits homophily, a greater number of link connections between similar nodes.¹ The network generator uses link density (ld) and homophily (dh) to govern network formation. A simplified version of the pseudo-code for this method is shown in Figure 1. For our model, we assumed predefined values for ld and dh . The nodes of the graph represent the individuals (agents) in the simulation, who can be considered as car drivers.

```

G = Null
repeat
  sample r from uniform distribution U(0, 1)
  if r ≤ ld then

    randomChooseSource(G)
    determineCandidateSink(dh, G)
    pickSink()      ▷ based on power-law distribution
    connect(source, sink)
  else
    add a new node to G
  end if
until desired number of nodes added to the network

```

Figure 1: Synthetic friendship network generator

We use a weighted voting approach (also known as a structure based method) to determine an agent’s decision with regard to its neighbors. The weight for each of an agent’s neighbors is computed using a normalized value of that neighbor’s centrality value as shown in Equation 1.

$$\text{weight}_{i,j} = \frac{C_j}{\sum_{k=1}^{Deg_i} C_k} \quad (1)$$

¹Commonly described as “birds of a feather flock together” [26]

This equation shows the weight of the link connecting neighbor j to node i . C refers to the corresponding centrality value (degree, betweenness and closeness). Also, Deg_i denotes the number of neighbors for node i . The top 10 percent of the population of agents with the greatest centrality values are assumed to be the key elements of a society. At the beginning of our experiments, all of the agents follow a single norm; in other words, all of them have learned (through social learning [32]) to always drive on one side of the road. In our implementation, each agent has a utility value defined for each of four possible cases: Up-Left, Up-Right, Down-Left and Down-Right, where Up and Down determine the section of road, and Left and Right determine the direction an agent drives. These values are updated while receiving payoffs based on the matrix shown in Table 1.

In our experiments, we compare the penetration of norm changing behaviors that emanate from key members of a society vs. other cases. We compare emanation from the top to emanation from the middle and bottom 10 percent of the population. At the beginning of the simulation, the agents (nodes) are ranked based on their centrality value to determine the top, middle and bottom agents. The utility value of these agents is kept fixed. Neighbors of these agents continue updating their behavior until a new norm emerges in the system. Figure 2, Figure 3, and Figure 4 show the number of iterations required for each case to converge. The population of agents contained 100 agents, and the reported results show the average values over 20 runs.

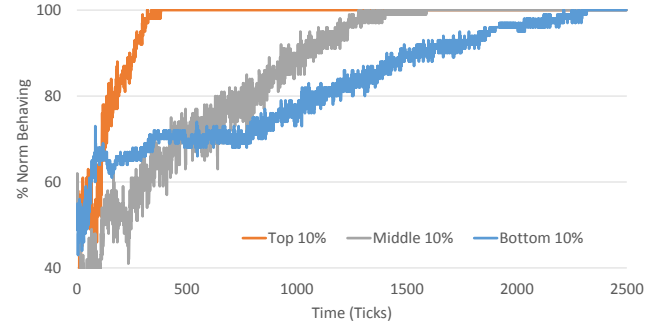


Figure 2: Average number of iterations until the emergence of one norm, when using degree centrality to determine key agents.

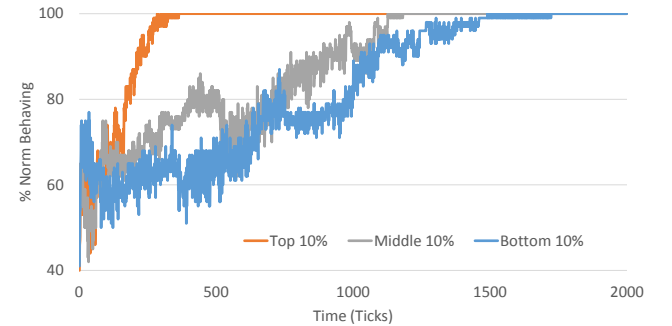


Figure 3: Average number of iterations until the emergence of one norm, when using betweenness centrality to determine key agents.

The pattern observed in all of three cases was very similar.

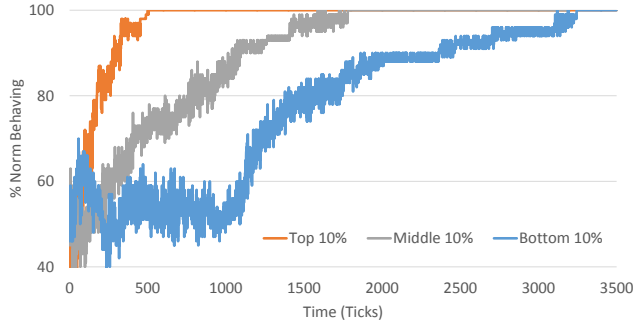


Figure 4: Average number of iterations until the emergence of one norm, when using closeness centrality to determine key agents.

When the norm propagation starts from the top 10% of the population, the norm emerges much faster compared to the other cases. Moreover, there is a fairly sizable difference among top, middle and bottom agents. The magnitude of difference between the top and middle 10% is more than the difference between the middle and bottom. These results are consistent with the role of connectors in tipping point theory.

6. STICKINESS FACTOR

According to the tipping point theory, the extent and rate of emerging social norms in a society is not only related to the members of the society, but also related to the content of the message. An effective message needs to be interesting or “sticky” enough to remain in agents’ minds. This factor is almost completely independent of the society and its structure, and is a property of the idea.

As Gladwell [21] points out, it is potentially very complicated to determine if a certain message has the necessary stickiness or not, but one characteristic that is usually common to sticky ideas is that it frequently returns to a person’s mind. This could be in the form of a desire to sit and watch a popular TV show every night, or in a more extreme case, a clinical addiction to smoking or gambling. Conventional marketing and advertising domains refer to this phenomenon as *rule of 27*. According to this rule, a message (advertisement) should be seen at least 27 times, if the message is going to stick [28].

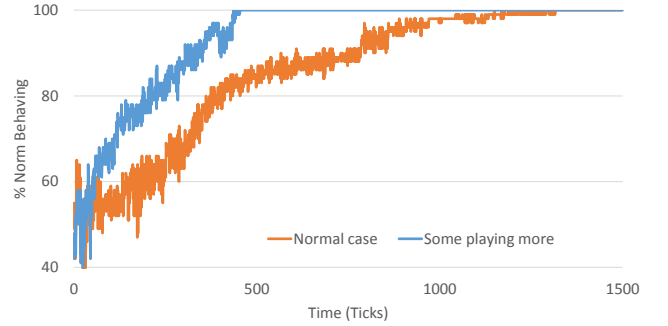


Figure 5: Average number of iterations until the emergence of one norm, when 2 out of 4 agents with fixed utility values play twice with each agent that they encounter.

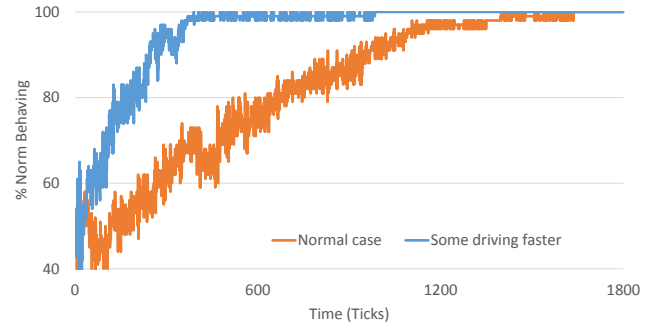


Figure 6: Average number of iterations until the emergence of one norm, when 2 out of 4 agents with fixed utility values go (drive) faster.

In order to model this property, we assume that the stickiness is represented by the number of games that an agent plays with another agent. Therefore a higher number of games will result in the same effect as a stickier belief. In our experiments, this idea is evaluated in two different ways. The first way is to increase the number of games that a certain set of agents play. The second way is to have a certain number of agents driving faster than other agents to be exposed to more cars.

Figures 5 and 6 show results related to these two cases. In both cases, we have our original 100 agents plus a group of 2 agents which have a fixed preference to drive on either the left or right. In the first scenario, one group of agents plays two games each time it encounters another agent. In the second scenario, one group of agents moves faster. Both of these scenarios lead to the same effect: increasing the number of times that an agent is exposed to an idea. This simulates the property of frequently returning to a person’s mind. In both cases, when the stickiness factor is implemented, the entire system converges to a single norm faster.

7. POWER OF CONTEXT

The third element of the tipping point theory refers to the power of context. As Gladwell points out: it is possible to be a better person on a clean street or in a clean subway, than in one littered with trash and graffiti [21]. The idea is mostly based on what known in criminology as the theory of *broken windows* [37]. According to this theory, slight changes in the environment could result in tipping effects over the whole society.

In order to apply this part of the tipping point theory, we use ideas from the methods for studying fads and cascading effects in networks [36]. First, we build a network using the same approach described in Section 5. Then, we assign a threshold value for each agent. Similar to the probabilistic information cascade models, if the cumulative value of the perceived cascade is less than the threshold, nothing will change. If it’s higher, the agent will change its current behavior, which in our scenario would result in driving on the other side of the road. Figure 7 shows the percentage of times that a norm emerged in the system for a set of threshold values. The columns show the average results over 20 runs. Agents were selected randomly as a source of a small initial shock in the network, which results in negating the current payoff values for driving on each side of the road. The frequency of shocks is determined randomly. The system runs until it reaches some fixed iteration number (50,000), unless a different norm is observed. This experiment illustrates how minor shocks can shape a population fad, resulting in a population-level behavior change. The shocks (pulses) in this model can be viewed as any of the small changes that tipping point theory predicts can result in large changes in the whole society. According to the results presented in Figure 7, thresholds as small as 0.02 can lead to the emergence of norms in the system in almost 5 percent of the experiments. The computed values for each agent are compared to its tipping point value (normalized between 0 and 1).

There is a second aspect to the power of context, which refers to the number of people in groups. The Rule of 150 says that the size of groups is a subtle contextual factor that makes a big difference. This number is referred as *Dunbar’s number* [18], after the anthropologist who originally proposed the idea. In groups with fewer than 150 members, people will cooperate relatively easily and rapidly become infected with the community ethos. Once that threshold is crossed people begin to behave very differently. 150 is our *social channel capacity* as determined on the basis of personal loyalties and 1-on-1 contacts. Beyond the tipping point of 150 the group dynamics simply become too complex. For the aver-

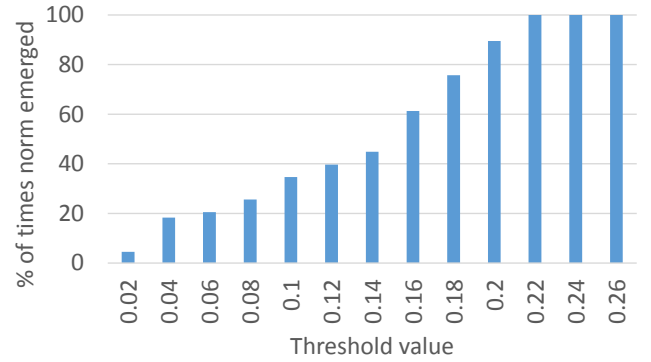


Figure 7: Percentage of times that a norm emerges in the population, when agents have different threshold values for activating.

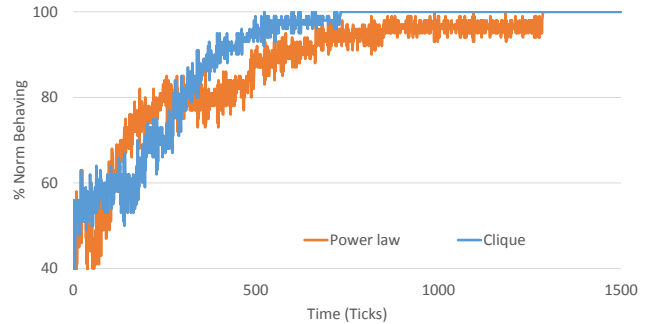


Figure 8: Average number of iterations until the emergence of one norm, when the network structure of agents follows a power-law distribution and when the network is a complete clique.

age person there are just too many relationships to manage. The group then becomes divided and alienated, and usually splits into two. Smaller groups have been shown to be more effective at tasks than larger groups. This may be due to biological limitations of humans which make it very difficult for them to handle a larger community.

With the growth of virtual social media sites and the spread of online groups, there has been renewed interest in evaluating the importance of this limit on Facebook [19], Myspace [20] and within massively multiplayer online role-playing games (MMORPGs). The pivotal issue here is that a person cannot maintain a close relationship with all of the members of a larger group which ultimately sabotages its success. Having a direct connection with each member of the group is a necessary component to having a positive social relationship.

We propose using a clique structure to illustrate this idea. In a clique each node has a direct edge to all of other nodes. There are $n * (n - 1)$ edges in the resulting graph. We opt to use a directed graph, as that seems to be the general assumption for friendship networks. We compare the emergence of driving norms in a network generated using the method described in Section 5. It should be noted that having more edges does not result in faster convergence. More connections makes the diffusion of ideas easier, while it makes it

harder for the agents to find an idea that all agents like. In a clique structure the major voting approach and the weighted voting approach (using the number of edges) are the effectively same, so neither of them elicits earlier norm emergence. Figure 8 shows the number of iterations that were required on average for the two cases to reach norm emergence. The driving norm emerged faster in case of the clique structure than in the power-law degree distribution network. This shows the potential benefit of such a structure in constructing agent systems, at least for ideal cases.

8. CONCLUSION AND FUTURE WORK

Norms are complex social behaviors that have been extensively studied in sociology, psychology, and other related fields. Most normative architectures draw upon theories from the social sciences. The theory of tipping points has inspired much research in different disciplines. For this paper, we model some of the well-known elements of this theory, as applied to networked agent populations. We illustrate how three of principle ideas including key few members, stickiness factor, and the role of environment can affect the process of norm emergence. This paper is meant to be an initial step for convincing the NorMAS community of the importance of tipping point theory concepts. For future work, we are interested in mapping the performance of our normative model to a real-dataset.

9. REFERENCES

- [1] T. Balke, T. Roberts, M. Xenitidou, and N. Gilbert. Modelling energy-consuming social practices as agents. In *Social Simulation Conference*, Barcelona, 2014.
- [2] R. Beheshti, A. M. Ali, and G. Sukthankar. Cognitive social learners: an architecture for modeling normative behavior. 2015.
- [3] R. Beheshti, M. Analui, and B. Minaei-Bidgoli. A pairwise classification method with support vector machines, based on game theory and alpha-beta algorithm. In *The 2nd International Conference of Iranian Operation Research Society*. Civilica, 2009.
- [4] R. Beheshti, R. Barmaki, and N. Mozayani. Negotiations in holonic multi-agent systems. In *The Seventh International Workshop on Agent-based Complex Automated Negotiations (ACAN2014)*, 2014.
- [5] R. Beheshti and N. Mozayani. Predicting opponents offers in multi-agent negotiations using artmap neural network. In *Future Information Technology and Management Engineering, 2009. FITME'09. Second International Conference on*, pages 600–603. IEEE, 2009.
- [6] R. Beheshti and N. Mozayani. A new mechanism for negotiations in multi-agent systems based on artmap artificial neural network. In *Agent and Multi-Agent Systems: Technologies and Applications*, pages 311–320. Springer, 2011.
- [7] R. Beheshti and N. Mozayani. Homan, a learning based negotiation method for holonic multi-agent systems. *Journal of Intelligent and Fuzzy Systems*, 26(2):655–666, 2014.
- [8] R. Beheshti, N. Mozayani, and A. T. Rahmani. A new hybrid evolutionary method with ant colony and pso algorithms based on fuzzy decision making. In *The 2nd International Conference of Iranian Operation Research Society*. Civilica, 2009.
- [9] R. Beheshti and A. Rahmani. A multi-objective genetic algorithm method to support multi-agent negotiations. In *Future Information Technology and Management Engineering, 2009. FITME'09. Second International Conference on*, pages 596–599. IEEE, 2009.
- [10] R. Beheshti and G. Sukthankar. Extracting agent-based models of human transportation patterns. In *Social Informatics (SocialInformatics), 2012 International Conference on*, pages 157–164. IEEE, 2012.
- [11] R. Beheshti and G. Sukthankar. An agent-based transportation simulation of the ucf campus. In *SwarmFest 2013: 17th Annual Meeting on Agent-Based Modeling & Simulation*, 2013.
- [12] R. Beheshti and G. Sukthankar. Analyzing agent-based models using category theory. In *Web Intelligence (WI) and Intelligent Agent Technologies (IAT), 2013 IEEE/WIC/ACM International Joint Conferences on*, volume 2, pages 280–286. IEEE, 2013.
- [13] R. Beheshti and G. Sukthankar. Improving markov chain monte carlo estimation with agent-based models. In *Social Computing, Behavioral-Cultural Modeling and Prediction*, pages 495–502. Springer, 2013.
- [14] R. Beheshti and G. Sukthankar. A hybrid modeling approach for parking and traffic prediction in urban

- simulations. *AI & SOCIETY*, pages 1–12, 2014.
- [15] R. Beheshti and G. Sukthankar. A normative agent-based model for predicting smoking cessation trends. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, pages 557–564. International Foundation for Autonomous Agents and Multiagent Systems, 2014.
- [16] G. Boella and L. van der Torre. An architecture of a normative system: counts-as conditionals, obligations and permissions. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*, pages 229–231. ACM Press, 2006.
- [17] R. Conte, G. Andrighetto, and M. Campennl. *Minding norms: Mechanisms and dynamics of social order in agent societies*. Oxford University Press, 2013.
- [18] R. I. Dunbar. Neocortex size as a constraint on group size in primates. *Journal of Human Evolution*, 22(6):469–493, 1992.
- [19] R. I. Dunbar. Primates on facebook. 2009. Retrieved from: <http://econ.st/1qgCrDL>.
- [20] J. W. Gibbons. *Modeling Content Lifespan in Online Social Networks Using Data Mining*. PhD thesis, University of Kansas, 2014.
- [21] M. Gladwell. *The tipping point: How little things can make a big difference*. Hachette Digital, Inc., 2006.
- [22] M. Granovetter. Threshold models of collective behavior. *American Journal of Sociology*, 83(6), 1978.
- [23] M. Grodzins. Metropolitan segregation. *Scientific American*, 1957.
- [24] C. D. Hollander and A. S. Wu. The current state of normative agent-based systems. *Journal of Artificial Societies and Social Simulation*, 14(2):6, 2011.
- [25] G. Malcolm. *The tipping point*. Little Brown, New York, 2000.
- [26] M. McPherson, L. Smith-Lovin, and J. M. Cook. Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, 27(1):415–444, 2001.
- [27] M. Neumann. Norm internalisation in human and artificial intelligence. *Journal of Artificial Societies and Social Simulation*, 13(1):12, 2010.
- [28] D. Pepper. *How to Start and Grow Your Lawn Care Maintenance Business*. Lulu Enterprises Incorporated, 2008.
- [29] B. T. R. Savarimuthu, S. Cranefield, M. K. Purvis, and M. A. Purvis. Norm emergence in agent societies formed by dynamically changing networks. *Web Intelligence and Agent Systems*, 7(3):223–232, 2009.
- [30] T. C. Schelling. *Micromotives and macrobehavior*. WW Norton & Company, 2006.
- [31] O. Sen and S. Sen. Effects of social network topology and options on norm emergence. In *Coordination, Organizations, Institutions and Norms in Agent Systems V*, pages 211–222. Springer, 2010.
- [32] S. Sen and S. Airiau. Emergence of norms through social learning. In *IJCAI*, volume 1507, page 1512, 2007.
- [33] D. Villatoro, G. Andrighetto, J. Sabater-Mir, and R. Conte. Dynamic sanctioning for robust and cost-efficient norm compliance. In *IJCAI*, volume 11, pages 414–419, 2011.
- [34] D. Villatoro, S. Sen, and J. Sabater-Mir. Topology and memory effect on convention emergence. In *Proceedings of the 2009 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology-Volume 02*, pages 233–240. IEEE Computer Society, 2009.
- [35] X. Wang, M. Maghami, and G. Sukthankar. Leveraging network properties for trust evaluation in multi-agent systems. In *Proceedings of the IEEE/WIC/ACM International Conferences on Web Intelligence and Intelligent Agent Technology-Volume 02*, pages 288–295. IEEE Computer Society, 2011.
- [36] D. J. Watts. A simple model of fads and cascading failures on sparse switching networks. In *Economics with Heterogeneous Interacting Agents*, pages 13–25. Springer, 2001.
- [37] J. Q. Wilson and G. L. Kelling. Broken windows. *Atlantic Monthly*, 249(3):29–38, 1982.
- [38] H. P. Young. The economics of convention. *The Journal of Economic Perspectives*, pages 105–122, 1996.
- [39] C. Yu, M. Zhang, F. Ren, and X. Luo. Emergence of social norms through collective learning in networked agent societies. In *Proceedings of the International Conference on Autonomous Agents and Multi-agent Systems*, pages 475–482. International Foundation for Autonomous Agents and Multiagent Systems, 2013.